

We claim:

1        1. A method for copying data, written by a host computer to a local data  
2 storage array, to a backup copy of the data on a remote storage system, wherein a  
3 local array controller, connected to the local data storage array, is coupled to a  
4 remote array controller connected to the remote storage system, the method  
5 comprising the steps of:

6                assigning a first current sequence number as the look-ahead limit for a write  
7 command associated with the data;

8                upon command completion, storing a second current sequence number as the  
9 write command sequence number;

10               placing the command in a queue;

11               sending the command sequence number and look-ahead limit, with the  
12 command associated therewith, to the remote array controller;

13               on reception of the received command by the remote array controller,

14               if there are any outstanding writes in progress, then

15                if the sequence number of the received command is less than a  
16 present look-ahead fence, then replacing the current look-ahead fence with the  
17 minimum of the current look-ahead fence and the look-ahead limit for the command;  
18 otherwise,

19                if the sequence number of the received command is not less  
20 than the current look-ahead fence, then putting the command in a waiting queue;

21                if there were no writes in progress, then setting the current look-  
22 ahead fence for the remote array controller to the look-ahead limit for the command;

23                executing the command at the remote controller;

24                when a write operation completes at the remote array controller, then

25                if the look-ahead limit for the command is equal to the current look-  
26 ahead fence, then replacing the current look-ahead fence with the minimum value of  
27 the look-ahead limit of all currently outstanding said commands; and

28                continuing to process said outstanding commands by directing each said  
29 command in the waiting queue back to the sending step.

1           2. The method of claim 1, wherein a sliding window protocol is used in  
2 said sending step to protect against out-of-order receptions at the remote site.

1           3. A method for copying data, written by a host computer to a local data  
2 storage array, to a remote storage system wherein a remote controller is coupled to  
3 a local array controller connected to the local data storage array, the method  
4 comprising the steps of:

5           establishing a current look-ahead fence equal to a maximum number of  
6 commands which are allowed to be outstanding between the local array controller  
7 and the remote array controller;

8           sending a command sequence number and look-ahead limit, with an  
9 associated write command, to the remote controller;

10           on reception of the received command by the remote controller,

11           if there are any outstanding writes in progress, then

12           if the sequence number of the received command is less than  
13 the present look-ahead fence, then replacing the current look-ahead fence with the  
14 minimum of the current look-ahead fence and the look-ahead limit for the command;  
15 otherwise,

16           if the sequence number of the received command is not less  
17 than the current look-ahead fence, then putting the command in a waiting queue;

18           if there are no writes in progress, then setting the current look-ahead  
19 fence for the remote controller to the look-ahead limit for the command;

20           executing the command at the remote controller; and

21           when the command completes at the remote controller, then

22           if the look-ahead limit for the command is equal to the current look-  
23 ahead fence, then replacing the current look-ahead fence with the minimum value of  
24 the look-ahead limit of all currently outstanding said commands.

1           4. The method of claim 3, wherein a sliding window protocol is used in  
2 said sending step to protect against out-of-order receptions at the remote site.

1           5. A method for copying data, written by a host computer to a local data  
2 storage array, to a backup copy of the data on a remote storage system, wherein a  
3 local array controller, connected to the local data storage array, is coupled to a  
4 remote array controller connected to the remote storage system, the method  
5 comprising the steps of:

6           sending a plurality of write commands, associated with the data, from the  
7 local array controller to the remote array controller;

8           in the situation wherein a first one of the write commands completes before a  
9 second one of the write commands starts at the local site, then inhibiting said data  
10 associated with the second one of the write commands from being recorded to media  
11 on the remote storage system before said data associated with the first one of the  
12 write commands is recorded to media on the remote storage system; and

13           using a sliding window protocol in said sending step to protect against out-  
14 of-order receptions at the remote site.

1           6. A method for copying data, written by a host computer to a local data  
2 storage array, to a backup copy of the data on a remote storage system, wherein a  
3 local array controller, connected to the local data storage array, is coupled to a  
4 remote array controller connected to the remote storage system, the method  
5 comprising the steps of:

6           sending a plurality of write commands, associated with the data, from the  
7 local array controller to the remote array controller; and

8           in the situation wherein a first one of the write commands completes before a  
9 second one of the write commands starts at the local site, then inhibiting said data  
10 associated with the second one of the write commands from being recorded to media  
11 on the remote storage system before said data associated with the first one of the  
12 write commands is recorded to media on the remote storage system.

1           7. The method of claim 6, wherein a sliding window protocol is used in  
2 said sending step to protect against out-of-order receptions at the remote site.

1           8. A method for copying data, written by a host computer to a local data  
2 storage array, to a backup copy of the data on a remote storage system, wherein a  
3 local array controller, connected to the local data storage array, is coupled to a  
4 remote array controller connected to the remote storage system, wherein a plurality  
5 of write commands, associated with the data, are sent from the local array controller  
6 to the remote array controller, the method comprising the step of:

7           in the situation wherein a first write command completes before a second  
8 write command starts at the local site, then inhibiting said data associated with the  
9 second write command from being recorded to media on the remote storage system  
10 before the first write command is recorded to media at the remote site.

1           9. The method of claim 8, wherein a sliding window protocol is used to  
2 protect against out-of-order receptions at the remote site.

1           10. A method for tuning of a link between a local array controller and a  
2 remote array controller to compensate for the distance therebetween, the method  
3 comprising the steps of:

4           establishing a connection between the local array controller and the remote  
5 array controller;

6           sending a plurality of test frames from the local array controller to the  
7 remote array controller;

8           sending a response from the remote array controller to the local array  
9 controller for each of the test frames received by the remote controller;

10           receiving each said response at the local array controller;

11           determining two-way transit times for time elapsed between the sending of  
12 each one of said test frames sent and the receiving thereof;

13           calculating an average value for the transit times; and

14           tuning the link by increasing the average value by a contingency factor to  
15 generate a timeout value for communications over the link.

1           11. The method of claim 10, wherein the contingency factor is  
2 approximately 100.

1           12. The method of claim 10, wherein the contingency factor is user-variable  
2    to compensate for factors other than distance.